

OK TO ENTER: /TP/

DOCKET NO.: MSFT-1742/301617.1
Application No.: 10/602,283
09/19/2008
Notice of Allowance Dated: April 21, 2008

PATENT

This listing of claims will replace all prior versions, and listings, of claims in the application.

Listing of Claims:

1. (Currently Amended) A method for resynchronizing multiple copies of a database after a divergence in transaction history, the database, D1, and its copies, D1', having database pages, logs, and log records, comprising:

identifying the a point when the divergence in transaction history occurred using the database and database copy logs, the point represented by a fail over log sequence number (FOLSN), wherein the divergence in transaction history comprises at least one of:

a condition in which D1 fails for a period of time during which D1' remains operational, and D1 is brought back online; or

a condition in which D1 has provided data to the database and the data has not been received by D1'; and

processing the database log records created after the divergence in transaction history to populate from D1 to D1' or from D1' to D1 changes and identified content, the log records having log sequence numbers for use in processing.

2. (Original) The method as recited in claim 1, further comprising setting a maximum page log sequence number (MPLSN) to equal the FOLSN.

3. (Original) The method as recited in claim 2, further comprising performing a scan of the log of either of D1 or D1' to obtain a first log record that occurs after the FOLSN until an end-of-log on either of D1 or D1' is reached.

4. (Original) The method as recited in claim 3, further comprising checking to determine if there is a log record to analyze.

5. (Original) The method as recited in claim 4, further comprising replaying log records on D1 starting with log records occurring before an identified last checkpoint before the FOLSN and continuing to the FOLSN.

DOCKET NO.: MSFT-1742/301617.1

PATENT

Application No.: 10/602,283

Notice of Allowance Dated: April 21, 2008

6. (Original) The method as recited in claim 5, further comprising removing the log records after the FOLSN from either of D1 or D1'.

7. (Original) The method as recited in claim 6, further comprising copying log records from the FOLSN to the MPLSN from D1 to D1' or from D1' to D1.

8. (Currently Amended) The method as recited in claim 7, further comprising performing the operations recorded from the FOLSN to the MPLSN on D1 or D1' pages.

9. (Original) The method as recited in claim 8, further comprising preventing the use of D1 or D1' until the operations of claim 8 are completed.

10. (Currently Amended) The method as recited in claim 9, further comprising clearing the a MPLSN value.

11. (Original) The method as recited in claim 10, further comprising recovering D1 and/or D1'.

12. (Original) The method as recited in claim 11, further comprising catching up D1 with D1' or vice versa.

13. (Currently Amended) The method as recited in claim 4, further comprising performing a check to determine determining if the log record occurring after the FOLSN is a page format log record.

14. (Currently Amended) The method as recited in claim 4, further comprising performing a check to determine determining if the log record occurring after the FOLSN is a page update log record.

DOCKET NO.: MSFT-1742/301617.1

Application No.: 10/602,283

Notice of Allowance Dated: April 21, 2008

PATENT

15. (Original) The method as recited in claim 14, further comprising comparing a previous page log sequence number (PPLSN) of the log record occurring after the FOLSN with the FOLSN.

16. (Currently Amended) The method as recited in claim 15, further comprising upon determining that the PPLSN is less than the FOLSN, retrieving the current contents of the D1 pages or D1' pages.

17. (Original) The method as recited in claim 16, further comprising determining if the D1 pages or D1' pages are de-allocated.

18. (Original) The method as recited in claim 17, further comprising sending a dummy page from D1 to D1' or from D1' to D1 with a page LSN (PLSN) field of the log record set in a current end-of-log (EOL) log sequence number on D1 or D1'.

19. (Currently Amended) The method as recited in claim 18, setting the a value of the MPLSN to the value of the PLSN if the PLSN is larger than the MPLSN.

20. (Currently Amended) The method of claims 13, 14, or 19, further comprising obtaining the a next log record.

21. (Original) A computer readable medium having computer readable instructions to instruct a computer to perform the method as recited in claim 1.

22. (Currently Amended) A system for the resynchronization of a database having multiple copies after the an occurrence of a divergence in transaction history comprising:

database log records, the database log records comprising log sequence numbers representative of database operations and transactions;

a means for identifying the a point of a the divergence in transaction history, operating on the database log records and log sequence numbers to identify the divergence in transaction history, wherein the divergence in transaction history comprises at least one of:

DOCKET NO.: MSFT-1742/301617.1

PATENT

Application No.: 10/602,283

Notice of Allowance Dated: April 21, 2008

a condition in which D1 fails for a period of time during which D1' remains operational, and D1 is brought back online; or

a condition in which D1 has provided data to the database and the data has not been received by D1'; and

a means for comparing information about the log records, the means for comparing using page sequence numbers and fail over sequence numbers to determine if changes have been made to the log records after the divergence in transaction history.

23. (Original) The system as recited in claim 22, wherein the means for identifying comprises a computing application.

24. (Original) The system as recited in claim 22, wherein the means for comparing information comprises a computing application.

25. (Currently Amended) A method for resynchronizing multiple copies of a database after a divergence in transaction history, the database, D1 operating on server computer S1, and its copies, D1' operating on server computers S2, having database pages, logs, and log records, comprising:

identifying ~~a~~ the divergence in transaction history, the divergence represented by a log sequence number (LSN), wherein the divergence in transaction history comprises at least one of:

a condition in which D1 fails for a period of time during which D1' remains operational, and D1 is brought back online; or

a condition in which D1 has provided data to the database and the data has not been received by D1';

on the server computer being recovered, scanning ~~the~~ a log file to identify ~~the~~ a set of pages to be replaced;

replacing those pages on the server computer being recovered with their images from the ~~the~~ a non-recovering server computer;

rolling forward ~~the~~ a transaction log from ~~the~~ a last checkpoint through a failover LSN on any of those pages being replaced on the server computer being recovered;

DOCKET NO.: MSFT-1742/301617.1

PATENT

Application No.: 10/602,283

Notice of Allowance Dated: April 21, 2008

copying the log and log records from the non-recovering server computer from the failover LSN to a maximum LSN; and
replaying the copied log records.

26. (Currently Amended) A method for resynchronizing multiple copies of a database after a divergence in transaction history, the database, D1, and its copies, D1', having database pages, logs, and log records, comprising:

identifying the a point when the divergence in transaction history occurred using the database and database copy logs, the instance divergence represented by a fail over log sequence number (FOLSN) , wherein the divergence in transaction history comprises at least one of:

a condition in which D1 fails for a period of time during which D1' remains operational, and D1 is brought back online; or

a condition in which D1 has provided data to the database and the data has not been received by D1';

creating a list of pages having been changed; and

processing the database log records created after the divergence in transaction history to populate from D1 to D1' or from D1' to D1 changes and identified content, the log records having log sequence numbers for use in processing.

27. (Original) The method as recited in claim 26, further comprising setting the MPLSN to equal the FOLSN.

28. (Original) The method as recited in claim 27, further comprising creating an in-memory hash table, the in-memory hash table having data identifying database pages which have changed.

29. (Currently Amended) The method as recited in claim 28, further comprising scanning the log of D1 or D1' to obtain the a first log record occurring after the FOLSN, wherein the scan is performed until an end-of-log (EOL) log sequence number is encountered.

DOCKET NO.: MSFT-1742/301617.1

PATENT

Application No.: 10/602,283

Notice of Allowance Dated: April 21, 2008

30. (Original) The method as recited in claim 29, further comprising performing a check to determine if there is a log record to process.

31. (Currently Amended) The method as recited in claim 30, further comprising upon determining there are no log records to process, looping over the pages found in the Changed Page Table (CPT) ~~CPT~~ and copying them to D1 or D1'.

32. (Original) The method as recited in claim 31, further comprising sending a dummy page and dummy ID by D1 to D1' or from D1' to D1, the dummy page having a page LSN (PLSN) being set to the EOL LSN on D1 or D1'.

33. (Original) The method as recited in claim 32, further comprising determining the MPLSN returned on pages from D1 or D1' and store it in the MPLSN.

34. (Currently Amended) The method as recited in claim 33, further comprising replaying the log records on D1 or D1' starting with ~~the~~ a last checkpoint before the FOLSN and continuing to the FOLSN.

35. (Original) The method as recited in claim 34, further comprising removing log records occurring after the FOLSN from D1 or D1'.

36. (Original) The method as recited in claim 35, further comprising copying the log records from the FOLSN to the MPLSN from D1 to D1' or D1' to D1.

37. (Original) The method as recited in claim 36, further comprising re-performing operations recorded from the FOLSN to the MPLSN on D1 or D1' pages.

38. (Original) The method as recited in claim 37, further comprising preventing the use of D1 or D1' until the operations of claim 37 are completed.

DOCKET NO.: MSFT-1742/301617.1
Application No.: 10/602,283
Notice of Allowance Dated: April 21, 2008

PATENT

39. (Currently Amended) The method as recited in claim 38, further comprising clearing ~~the a~~ MPLSN value.

40. (Original) The method as recited in claim 39, further comprising recovering D1 and/or D1'.

41. (Original) The method as recited in claim 40, further comprising performing log shipping between D1 and D1'.

42. (Currently Amended) The method as recited in claim 30, further comprising ~~performing a check to determine determining~~ if the log record is a page format record.

43. (Currently Amended) The method as recited in claim 42, further comprising ~~performing a check to determine determining~~ if the log record is a page update log record.

44. (Currently Amended) The method as recited in claim 43, further comprising ~~performing a check to determine determining~~ if the page is in the Changed Page Table (CPT) CPT.

45. (Currently Amended) The method as recited in claim 44, further comprising placing a page ID in the Changed Page Table (CPT) CPT.

46. (Currently Amended) The method as recited in claim 30, 43, or 44, further comprising obtaining ~~the a~~ next log record.